

=> d his

(FILE 'HOME' ENTERED AT 10:29:54 ON 17 OCT 2001)

FILE 'USPATFULL' ENTERED AT 10:30:35 ON 17 OCT 2001

L1 978 S TRANSLATION LOOKASIDE BUFFER#
L2 272 S TLB ENTRY
L3 195 S TLB ENTRIES
L4 304 S L2 OR L3
L5 184023 S VALID? OR STATUS
L6 79 S L5 (3A) L4
L7 3923 S CACHE (3A) MISS?
L8 3435 S CACHE (3A) HIT?
L9 2526 S L7 AND L8
L10 10 S L9 AND L6
L11 136 S (TLB (3A) (ENTR### OR LINE#)) (P) (CACHE (3A) (ENTR### OR
LIN
L12 120 S L5 AND L11
L13 69 S L9 AND L12
L14 7627 S MULTIPROCESS?
L15 40724 S PLURAL? (3A) PROCESS?
L16 44943 S L14 OR L15
L17 39 S L13 AND L16
SAVE ALL C09212291/L
L18 95 S TLB# (4A) SIZE?
L19 57 S TLB# (2A) SIZE?
L20 11 S L13 AND L18
L21 1 S TLB# SIZE# (4A) SAME
L22 8 S TLB# SIZE#

=> s (TLB (2W) (ENTR? OR LINE#)) (3W) (ASSOCIAT? OR CORRESPOND?)

1777 TLB
398658 ENTR?
1772819 LINE#
1096489 ASSOCIAT?
1603741 CORRESPOND?
L23 54 (TLB (2W) (ENTR? OR LINE#)) (3W) (ASSOCIAT? OR CORRESPOND?)

37 ANSWER 46 OF 54 USPATFULL

PI US 4797814 19890110

SUMM In prior U.S. Pat. No. 4,464,712, the page-translating **TLB**

entries correspond to page-size lines in the L2 cache,
which has a L2 cache directory separate from the TLB (i.e. DLAT).
Absolute-addresses outputted by the TLB upon each TLB entry replacement
operation locate and control the settings of replacement-candidate flag
bits R in the L2 entries to control the LRU replacement selection of
line entries in the L2 cache directory. This requires a TLB/L2
relationship in which the L2 cache has an L2 line size equal to the TLB
controlled page size (e.g. 4096 bytes).

L13 ANSWER 58 OF 69 USPATFULL

PI US 5317711 19940531

SUMM In a computer system with a cache memory, when a memory word is needed, the central processing unit (CPU) looks into the cache memory for a copy

of the memory word. If the memory word is found in the **cache** memory, a **cache "hit"** is said to have occurred, and the main memory is not accessed. Thus, a **FIGURE** of merit which can be used to measure the effectiveness of the **cache** memory is the "**hit**" ratio. The hit ratio is the percentage of total memory references in which the desired datum is found in the cache memory without accessing the main memory. When the desired datum is not found in the **cache** memory, a "**cache miss**" is said to have occurred and the main memory is then accessed for the desired datum. In addition, in many computer systems, there are

portions of the address space which are not mapped to the cache memory. These portions of the address space are said to be "uncached" or "uncacheable". For example, the addresses assigned to input/output

(I/O) devices are almost always uncached. Both a **cache miss** or an uncacheable memory reference result in an access to the main memory.

SUMM In the course of developing or debugging a computer system, it is often necessary to monitor program execution by the CPU or to interrupt one instruction stream to direct the CPU to execute certain alternate instructions. For example, a technique for testing a microprocessor in

a system under development uses an in-circuit emulator (ICE) which monitors the instruction stream and, where appropriate, takes control

of the microprocessor by forcing the microprocessor to execute an alternative program. To achieve this end, the ICE monitors the signals on the microprocessor's pins. When a monitored instruction in the program execution is encountered, the ICE causes alternative instructions to be executed for such purpose as reading or altering the internal state of the CPU. Hence, when the cache memory is implemented off-chip, the transactions between the cache memory and the CPU can be monitored by the ICE via the microprocessor's pins on the off-chip bus between the cache memory and the CPU. However, when the cache memory is implemented on-chip, the transactions between the cache and the CPU occur on an internal on-chip bus, which cannot be probed from the pins of the integrated circuit. As a result, debugging operations using an ICE in a system with an on-chip cache memory can be very difficult.

When the on-chip cache memory achieves a high hit ratio, only the relatively infrequent accesses to main memory due to **cache misses** or references to uncacheable portions of memory can be monitored from the pins.

DETD FIG. 2 is a block diagram showing the addressing scheme used in the instruction cache 102a of the cache system 102, which is shown in FIGS. 1a and 1b. As shown in FIG. 2, the higher order 20 bits of a virtual address (generated by CPU core 103, as shown in FIG. 1b), which is represented by block 202, is provided to the cache addressing mechanism represented by block 201. The remaining 10 bits of the memory word address are common between the virtual and the physical addresses. (The

lowest two bits are byte addresses, which are not used in cache addressing.) These ten common bits are directly provided to index into the cache memory 102a, represented by blocks 204 and 205. Block 205 represents the data portion of the cache line, which comprises four 32-bit memory words in this embodiment. Block 204 represents the "tag" portion of the cache data word; this tag portion contains both a "valid" bit and the higher order 20 bits of the memory word addresses of the data words stored in the cache line. (Since the addresses of memory words within the cache line are contiguous, the higher order 20 bits are common to all of the memory words in the cache line). The **valid** bit indicates that the cache word contains **valid** data. Invalid data may exist if the data in the cache does not contain a current memory word. This condition may arise, for example, after a reset period.

DETD Each virtual address is associated with a particular process identified by a unique "process id" PID, which is represented by block 203. Block 201 represents the virtual address to the physical address translation, which is performed using the TLB unit 103b-3 when the TLB is present. (FIG. 1b.) When the TLB is present, a TLB miss occurs if either a mapping between the virtual address and the corresponding physical address cannot be found in the 64 **entries** of the **TLB** unit 103b-3, the PID stored in the TLB unit 103b-3 does not match the PID of the virtual address, or if the **valid** bit in the data word is not set. Block 207 represents the determination of whether a

TLB miss has occurred. The TLB miss condition raises an exception condition, which is handled by CPU core 103. If a virtual address to physical address mapping is found, the higher order 20 bits of the physical memory word address is compared (block 206) with the memory address portion of the tag. The **valid** bit is examined to ensure the data portion of the **cache line** contains **valid** data. If the comparison (block 206) indicates a **cache hit**, the selected 32-bit word in the **cache line** is the desired data.

DETD If a **cache miss** is indicated, BIU 106 is invoked and CPU core 103 stalls until BIU 106 indicates that the requested data is available. A **cache miss** can also be generated when the memory access is to a "uncacheable" portion of memory. When BIU 106 receives a datum from main memory, the CPU core 103 executes either a "refill", a "fix-up", or a "stream" cycle. In a refill cycle, an instruction datum received (in the read buffer 106-3) is brought into the cache 102a. In a fix-up cycle, the CPU core 103 transitions from a refill cycle to execute the instruction brought out of the read buffer 106-3. In a stream cycle, the CPU core 103 simultaneously refills cache memory 102a and executes the instruction brought out of the read buffer 106-3. For uncacheable references, the CPU core 103 executes a fixup cycle to bring out the fetched memory word from the read buffer 106-3, but the uncacheable memory word is not brought into the cache memory 102a. Otherwise, the CPU core 103 executes refill cycles until the miss address is reached. At that time, a fixup cycle is executed. Subsequent cycles are stream cycles until the end of the 4-memory word block is reached and normal run operation resumes. If sequential execution is interrupted, e.g. a successful branch condition, refill cycles are executed to refill the cache before execution is resumed at the branch address.

DETD In this embodiment, processor 101 is further provided with five "reserved pins" RSVD[4:0], for receiving five signals used for testing purposes. FIG. 4 shows the signal on each of the pins RSVD[4:0] being received by input buffers 401, which is decoded by decoding logic 402 into five signals 403-407, respectively labelled CPU.sub.-- TRIB, CA.sub.-- TEST, ICE.sub.-- REQ, ICA.sub.-- INVB, and ICE.sub.-- ADDR.

In this embodiment, when the signal on the RSVD[2] pin is logic high, and the signals on the RSVD[3:4] pins are at logic low, the processor 101

is

is said to be in "debug" mode. In debug mode, the ICE.sub.-- ADDR signal asserted. Furthermore, the ICE.sub.-- INVB signal is asserted if the signal on the RSVD[1] pin is also asserted during debug mode. The ICE.sub.-- ADDR signal on terminal 407, which is asynchronous, indicates that echoing the signals on internal bus 108 between the CPU core 103a and the caches 102 is desired. The ICE.sub.-- INVB signal on terminal 406, which is synchronized with respect to clock signals SysClk, SysOut1 and SysOut2 and their respective complementary signals, indicates that a forced **cache miss** is desired. A forced **cache miss** is used to allow an external testing device, such as an ICE, to "jam" an instruction into the CPU core. The forced **cache miss** mechanism is described on copending Application entitled "Hardware Control Method for Monitoring an On-chip Internal Cache in a Microprocessor", by P. Bourekas et al, U.S. Ser. No. 07/715,525 filed Jun. 14, 1991 and which is hereby incorporated by reference in its entirety. The ability to monitor internal bus signals and the ability to force execution of alternative instructions are important features necessary to support the use of testing and debugging devices, such as an ICE.

DETD When the ICE.sub.-- ADDR signal is asserted, the signals on internal bus 108 are echoed on the output bus 153a (FIG. 1a), when the address/data bus 153a is idle. The address/data bus 153a is idle when neither read (e.g. a **cache miss**, or an uncacheable reference), write or DMA operations to the main memory are using the address/data bus 153a. In order to prevent the echoed signals on the internal bus 108 from being mistaken by the main memory as a request for read or write, the ALE (Address latch enable) signal of the Rd/Wr control bus 153b-1 (FIG. 1b) is disabled. Thus, to latch the signals on address/data 153a, an external circuit should use the rising edge of the clock signal SysClk.

DETD From the description of the run cycle above in conjunction with FIGS. 2 and 3, the person of ordinary skill in the art will appreciate that in the instruction phase (the second phase) of a run cycle, the least significant 10 bits of the memory address of the next instruction are used to address the instruction cache. The 8 higher bits of these 10 bits locate the tag. In the first phase of the next run cycle, a 22-bit tag is returned by the instruction cache. As discussed above, the 22-bit tag comprises the higher 20 bits shared by the addresses of the four memory words stored in the cache line, are lower address bit, and a "**valid**" bit. Hence, a memory word address corresponding to a **cache "hit"** can be constructed from concatenating the least significant 10 bits, which are used to address the instruction cache, and the 20-bit address portion of the tag returned. Therefore, FIG. 5 shows capturing 10 bits of the lower address bits Addr[11:2] (represented by bus 504 a in FIG. 5) on internal bus 108 half clock cycle prior to capturing the 20-bit address portion (represented by bus 504b in FIG. 5) of the tag returned. FIG. 5 shows that the timing for capturing the lower 10 bits Addr[11:2] of the address can be achieved by synchronizing the data in the address phase of the run cycle to clock signals SysOut1 and SysOut2. For output purpose, the 20-bit tag (Addr[31:12]) is output on bus 153a by the pins AD[31:12], bits Addr[11:4] are output on bus 153a by the pins AD[11:4]. Addr[3:2] are output on control signal pins Addr[3:2] of Rd/Wr control bus 153b-1.

DETD The address display mode is intended to allow gross, rather than fine instruction trace by a testing device, such as an ICE. For example, branch instruction executed while bus 153a is engaged in a DMA

transaction will not be traceable. Additionally, data echoed may not be valid when CPU core 153a stalls, such as when a TLB miss occurs.